

# Design and Operational Experience of a Microwave Cavity Axion Detector for the 20 – 100 $\mu\text{eV}$ Range

S. Al Kenany<sup>a</sup>, M.A. Anil<sup>b</sup>, K.M. Backes<sup>a</sup>, B.M. Brubaker<sup>c</sup>, S.B. Cahn<sup>c</sup>, G. Carosi<sup>d</sup>, Y.V. Gurevich<sup>c</sup>, W.F. Kindel<sup>b</sup>, S.K. Lamoreaux<sup>c</sup>, K.W. Lehnert<sup>b</sup>, S.M. Lewis<sup>a</sup>, M. Malnou<sup>b</sup>, D.A. Palken<sup>b</sup>, N.M. Rapis<sup>a</sup>, J.R. Root<sup>a</sup>, M. Simanovskaia<sup>a,\*</sup>, T.M. Shokair<sup>a</sup>, I. Urdinaran<sup>a</sup>, K.A. van Bibber<sup>a</sup>, L. Zhong<sup>c</sup>

<sup>a</sup>*Department of Nuclear Engineering, University of California Berkeley, Berkeley CA, 94720 USA*

<sup>b</sup>*JILA and the Department of Physics, University of Colorado and National Institute of Standards and Technology, Boulder CO, 80309 USA*

<sup>c</sup>*Department of Physics, Yale University, New Haven CT, 06511 USA*

<sup>d</sup>*Physical and Life Sciences Directorate, Lawrence Livermore National Laboratory, Livermore CA, 94551 USA*

## Abstract

We describe a dark matter axion detector designed, constructed, and operated both as an innovation platform for new cavity and amplifier technologies and as a data pathfinder in the 5 – 25 GHz range ( $\sim 20 - 100 \mu\text{eV}$ ). The platform is small but flexible to facilitate the development of new microwave cavity and amplifier concepts in an operational environment. The experiment has recently completed its first data production; it is the first microwave cavity axion search to deploy a Josephson parametric amplifier and a dilution refrigerator to achieve near-quantum limited performance.

**Keywords:** axion, dark matter, Josephson Parametric Amplifier, microwave cavity, standard quantum limit, superconducting magnet

## 1. Introduction

The axion is a hypothetical pseudoscalar arising from the Peccei-Quinn mechanism to protect the strong interaction from CP-violating effects. Also, an axion in the 1 – 100  $\mu\text{eV}$  mass range is a compelling dark matter candidate. A comprehensive review of the particle physics of the axion, its cosmological and astrophysical significance, and experimental searches for it can be found in Ref. [1].

The axion, like the  $\pi^0$ , can couple to two photons, one of which may be virtual. Pierre Sikivie thus proposed a practical detection strategy based on the resonant conversion of the dark matter axion to a single microwave photon carrying its full energy (mass + kinetic) in a high-Q cavity permeated by a strong magnetic field [2, 3]. The resonant conversion condition is that the axion mass  $m_a$  is close to the resonant frequency  $\nu_c$  of a cavity mode with an appropriate spatial profile; more precisely,  $|\delta\nu| = |\nu_c - m_a c^2/h| \lesssim \Delta\nu_c$ , where  $\Delta\nu_c$  is the linewidth of the mode. The axion kinetic energy distribution is generally assumed to be Maxwellian, parametrized by the virial velocity  $\langle v^2 \rangle^{1/2} \sim 10^{-3} c$ . The signal is thus monochromatic to a part in  $10^6$ , with linewidth  $\Delta\nu_a = m_a \langle v^2 \rangle/h \ll \Delta\nu_c$ ; see Fig. 1.

The axion conversion power in such a microwave cavity

detector is

$$P_{\text{sig}} = \left( g_\gamma^2 \frac{\alpha^2}{\pi^2} \frac{\hbar^3 c^3 \rho_a}{\Lambda^4} \right) \times \left( \frac{\beta}{1 + \beta} \omega_c \frac{1}{\mu_0} B_0^2 V C_{mnl} Q_L \frac{1}{1 + (2\delta\nu/\Delta\nu_c)^2} \right). \quad (1)$$

Eq. (1) is valid in any self-consistent set of units; the two sets of parentheses contain theory and detector parameters, respectively. On the theoretical side,  $\alpha$  is the fine-structure constant,  $\Lambda = 78 \text{ MeV}$  encodes the dependence of the axion mass on hadronic physics, and the local axion dark matter density  $\rho_a$  and coupling constant  $g_\gamma$  are the parameters that experiment can constrain. It is conventional to fix  $\rho_a = 0.45 \text{ GeV}/\text{cm}^3$  [4] and quote limits on  $g_\gamma$ , a model-dependent dimensionless coupling that is related to the physical coupling  $g_{a\gamma\gamma}$  appearing in the axion-photon Lagrangian by  $g_{a\gamma\gamma} = (g_\gamma \alpha / \pi \Lambda^2) m_a$ . In the KSVZ (DFSZ) benchmark axion model,  $g_\gamma = -0.97$  (0.36), independent of the axion mass.

The experimental parameters in Eq. (1) include the magnetic field strength  $B_0$ , the cavity volume  $V$ , and several factors characterizing the cavity mode, with  $\omega_c = 2\pi\nu_c$ . The coupling between the cavity mode and the receiver used to detect the signal, parametrized by  $\beta$ , reduces the quality factor from  $Q_0$  to  $Q_L = Q_0 / (1 + \beta)$ . The form factor  $C_{mnl}$  parametrizes the overlap between the cavity mode and the external magnetic field. For the cylindrical

\*Corresponding author

Email address: simanovskaia@berkeley.edu (M. Simanovskaia)

geometry commonly used in cavity axion detectors,  $C_{mnl}$  may be written

$$C_{mnl} = \frac{(\int d^3\mathbf{x} \hat{\mathbf{z}} \cdot \mathbf{e}_{mnl}^*(\mathbf{x}))^2}{V \int d^3\mathbf{x} \epsilon(\mathbf{x}) |\mathbf{e}_{mnl}(\mathbf{x})|^2}, \quad (2)$$

where  $\mathbf{e}_{mnl}(\mathbf{x})$  is the normalized electric field profile of the mode,  $\epsilon(\mathbf{x}) = 1$  is the dielectric constant inside the cavity, and  $B_0$  is axial and homogeneous. Nodes in the electric field profile lead to cancellations in the form factor, and thus it is only appreciable for low-order  $\text{TM}_{0n0}$  modes.

Inserting typical values for the detector described in this work, which are close to the limits of present technology, the peak signal power is  $P_{\text{sig}} \sim 5 \times 10^{-24}$  W at KSVZ coupling, a factor of  $10^9$  below room-temperature thermal noise power in a typical cavity bandwidth. Thus cryogenic operation and a low-noise receiver are necessary for any practical realization of a cavity axion detector. The intrinsic spectral resolution of the linear receivers used in cavity axion searches to date implies that the relevant noise bandwidth is  $\Delta\nu_a$  rather than  $\Delta\nu_c$ , and moreover that a measurement at any given cavity frequency simultaneously probes  $\sim \Delta\nu_c/\Delta\nu_a$  independent values of the axion mass. Further improvement in the signal-to-noise ratio (SNR) may then be obtained by averaging the cavity noise for a time  $\tau$ . The preceding discussion is formalized in the Dicke radiometer equation [5], in which the SNR is defined as

$$\Sigma = \frac{P_{\text{sig}}}{k_B T_{\text{sys}}} \sqrt{\frac{\tau}{\Delta\nu_a}}. \quad (3)$$

For any phase-insensitive linear receiver the system noise temperature  $T_{\text{sys}}$  may be written

$$k_B T_{\text{sys}} = h\nu N_{\text{sys}} = h\nu \left( \frac{1}{e^{h\nu/k_B T} - 1} + \frac{1}{2} + N_A \right), \quad (4)$$

where the three additive contributions correspond respectively to a blackbody gas in equilibrium with the cavity at temperature  $T$ , the zero-point fluctuations of the blackbody gas, and the input-referred added noise of the receiver. The price we pay for spectral resolution is the quantum limit  $N_A \geq 1/2$  on the added noise of any phase-insensitive linear receiver [6], which together with the second term in Eq. (4) implies a “standard quantum limit”  $N_{\text{sys}} \geq 1$  for microwave cavity axion detection. These limits imply that units of quanta, used throughout this work, are more appropriate than temperature units for sufficiently low-noise receivers.

The axion mass is a priori unknown, so the cavity must be tuned in small discrete steps after each measurement interval  $\tau$ ; the SNR for any given value of the axion mass is then given by a quadrature sum of terms with the form of Eq. (3). An average scan rate is obtained by solving for  $1/\tau$  and multiplying by the frequency step size. With the simplifying assumption that the detector parameters

remain approximately constant over its tuning range, the scan rate is

$$\frac{d\nu}{dt} \approx \frac{4}{5} \frac{Q_L Q_a}{\Sigma^2} \left( g_\gamma^2 \frac{\alpha^2}{\pi^2} \frac{\hbar^3 c^3 \rho_a}{\Lambda^4} \right)^2 \times \left( \frac{1}{\hbar \mu_0} \frac{\beta}{1 + \beta} B_0^2 V C_{mnl} \frac{1}{N_{\text{sys}}} \right)^2, \quad (5)$$

where  $Q_a = (\langle v^2 \rangle / c^2)^{-1}$  is the “quality factor” of the axion signal, and the factor of  $4/5$  comes from a sum over the squared Lorentzian factors characterizing the effect of the axion’s changing detuning from the cavity mode. This is a good approximation for any frequency step size  $\lesssim \Delta\nu_c/2$ ; qualitatively, a smaller step size implies that the same sensitivity can be achieved with smaller  $\tau$ , and the two effects cancel out in the scan rate. Note also that, while Eq. (1) is maximized at critical coupling ( $\beta = 1$ ) for a given  $Q_0$ , the scan rate is maximized for an overcoupled cavity with  $\beta = 2$ . Eq. (5) is the most useful figure of merit for the cavity axion search.




Figure 1: Schematic of the microwave cavity search for dark matter axions. The axion signal is designated by the narrow peak (red) within the bandpass of the cavity (pink).

The detector described in this work builds upon experience from the Axion Dark Matter eXperiment (ADMX), a larger platform designed to operate initially in the sub-GHz range. ADMX entered into operation in 1996 with a physical temperature of  $T \approx 1.5$  K [4]. The readout chain, initially led by a high electron mobility transistor amplifier (HEMT) [4], was later upgraded to include a microstrip-coupled SQUID amplifier (MSA) [7]. A dilution refrigerator has recently been incorporated into the setup, and commissioning is underway.

Our detector was conceived as both an innovation platform and a data pathfinder for the microwave cavity axion search in the 5 – 25 GHz range ( $\sim 20 - 100$   $\mu\text{eV}$  axion masses;  $1 \text{ GHz} = 4.136 \mu\text{eV}$ ). The next decade in frequency brings new challenges for both microwave cavities and receivers. On the cavity side, preserving the aspect ratio, the volume of the cavity  $V \sim \nu^{-3}$ , resulting in an

increasing penalty in conversion power with increasing frequency. This requires insightful cavity designs to maintain useful volume at higher frequencies without sacrificing the form factor or incurring an unacceptable density of intruder TE modes in the spectrum. The absolute machining and alignment tolerances will also become tighter for smaller structures, to avoid mode localization. Another challenge is that MSAs do not work well above a gigahertz; to circumvent this problem we have introduced Josephson parametric amplifiers (JPA) that are ideal for the 5 GHz range. Together with a dilution refrigerator (DR), built into our design from the beginning, the JPA allows us to push down towards the Standard Quantum Limit and thus remain competitive with lower-frequency limits despite the loss of sensitive volume.

At the same time, this frequency range presents an opportunity to deploy entirely novel technologies that promise dramatic new capability for the microwave cavity search. A technology that will be investigated in the near future is a receiver based on squeezed states of the vacuum to circumvent the Standard Quantum Limit. Innovations to be tested in the microwave cavity domain include photonic band gap resonators, designed to eliminate the spectrum of TE modes which otherwise mix with the  $TM_{010}$  mode of interest as it is tuned, resulting in a loss of sensitivity at mode crossings and making it more difficult to track the  $TM_{010}$  mode. Other schemes to be investigated are the application of distributed Bragg reflector based resonator schemes, and thin-film Type II superconducting coatings to improve the quality factor  $Q$  of the cavity.

## 2. Description of Experimental Setup

Fig. 2 presents an overview of the experiment, sited at the Wright Laboratory of Yale University, and its integration. The microwave cavity and the magnetically-shielded canister housing the JPA are assembled on a gantry suspended from the dilution refrigerator. The gantry assembly is lowered by crane into the bore of the magnet, which is located in a room below the floor level of the lab containing the electronics and computer control.

The DR precooling system and the magnet cryogenic system use pulse tube and Gifford-McMahon cryocoolers, respectively. This allows operation of our detector without external liquid cryogens, providing major simplification of operations, and a substantial reduction in operating costs. However, such a cryogen-free system relies on an uninterrupted supply of electrical power as the consequences of a magnet quench can be severe (see section 2.2).

The individual systems are described in detail below.

### 2.1. Dilution Refrigerator

The DR was manufactured in 2007 by VeriCold Technologies (subsequently acquired by Oxford Instruments), and was among the first commercially available external-cryogen-free DRs. The only major modification was the

replacement of the VeriCold pulse tube with a Sumitomo RDK-415D Gifford-McMahon cryocooler (in retrospect this was a poor choice because of excessive vibration). In addition, instead of using the DR's Lakeshore temperature bridge to control the mixing chamber plate, we have incorporated a custom temperature PID-feedback controller at the mixing chamber level, resulting in reduced delay and improved performance.


The cooling power of the DR is about  $150 \mu\text{W}$  at 127 mK, which is the system operating temperature. This operating temperature was chosen because at lower temperature the JPA system was extremely sensitive to mechanical vibrations. In addition, operating at higher temperature results in decreased temperature excursion, due to the increase in specific heat of the Cu mixing chamber plate and other metal components, from heat generated with mechanical motions and operation of the microwave switch. The cavity, JPA, mixing chamber, and still temperatures are monitored with calibrated ruthenium oxide sensors.

The support gantry is a tripod with copper alloy legs and copper rings at both ends, which are clamped to the bottom of the DR's mixing chamber plate and to the upper endcap of the cavity, respectively. The equilibration time for changes in the mixing chamber temperature to propagate to the cavity temperature (measured at the lower endcap) is on the order of a few minutes. This might be surprising as the cavity is constructed of stainless steel, which has very low thermal conductivity. However, the cavity has a thick Cu plating (0.125 mm) which provides a sufficient thermal link for a short equilibration time.



The cavity and gantry are thermally shielded with an extension of the DR still shield. This extension comprises two demountable sections with dimensions 17.8 cm OD  $\times$  45 cm and 13.0 cm OD  $\times$  53.3 cm for the upper and lower sections, respectively. At the bottom of the still shield extension, there is a G-10 fiberglass disk attached via a stainless steel rod, which centers the shield extension and maintains the 0.5 cm gap between the still shield and magnet bore. The still shield extension was originally constructed from 0.16 cm sheet copper with welded seams, but that was deformed when we experienced a magnet quench (more information in section 2.2). Subsequently, we constructed a replacement using heavily plated stainless steel (0.125 – 0.250 mm), prompted by our experience that a heavy copper plating provides a good thermal link.

The DR and magnet share a single common insulating vacuum space. Extensions of the 4 K and 77 K DR stages overlap the corresponding stages at the top of the magnet, and thermal contact is provided by finger stock that presses between the extensions and the corresponding magnet thermal stages and also by blackbody radiation between the overlapping surfaces. This design allows the DR/cavity/still shield assembly to be easily inserted into and removed from the magnet by use of an overhead hoist. The gas flow and electronic control lines are long enough that only a few cables must be disconnected to insert or remove the cryostat. The DR itself is supported

Josephson Parametric Amplifier



Microwave Cavity (copper)

 $^3\text{He}/^4\text{He}$  Dilution Refrigerator

9.4 Tesla, 10 Liter Magnet



Figure 2: Overview of the experiment and its integration.

from above, and the magnet is on a cart equipped with jack screws. When the DR is lifted, the magnet can be rolled away and the DR lowered again to its support, allowing easy and safe access to the assembly below the DR. To insert the DR, the magnet is rolled to the appropriate location, the DR and extension are lowered into the magnet, and the jack screws are used to raise the magnet until the corresponding DR and magnet cryostat vacuum flanges contact, after which the flanges are bolted together.

The DR operating alone has a cooldown time of about 14 hours; with the gantry, cavity, and magnet, the cooldown time to 127 mK cavity temperature is 72 hours.

## 2.2. Magnet

The magnet and its controller were designed and manufactured by Cryomagnetics, Inc., based on an existing design for a gyrotron system with similar field requirements. As mentioned already, the magnet is a pulse-tube cooled system, requiring no external cryogenics or helium gas, and operates in a persistent mode. The magnetic field is ramped up to a maximum of 9 T over 8 hours; the rate is limited by heating of the magnet itself and by eddy current heating of the cavity/gantry/shield assembly. An interesting feature of the heating due to the ramping is that the heat load inferred from the rise in cavity temperature is largest when the field is between about 4 and

7 T, which we attribute to electrons being polarized in the stainless steel cavity body.

Our design specification was that the field needs to be homogenous at a level such that the perpendicular (radial) field at the inner surface of the cylindrical cavity barrel is less than 50 G. For future versions of this axion detector, we are interested in coating this inner surface with superconducting thin films to improve  $Q_0$ . The mentioned design specification is a requirement for a thin film superconducting layer to remain effective in the presence of a strong magnetic field. The idea is that the vortices formed by the perpendicular component of the field would be of sufficiently low density and mobility to not contribute significantly to radiofrequency power loss in the film [8]. The modification of the gyrotron design includes a set of superconducting coils in series with the main magnet that cancel the axial field over a 15.2 cm region to less than 50 G. This is the location of the JPA and its shield.

The magnet bore diameter is 14 cm inside the main solenoid and steps up to 18.4 cm further up, where the bucking coils are supported. The effective and homogeneous region in the field maximum extends axially over at least 25 cm.

An unfortunate disadvantage of a cryogen free system is that it relies on an uninterrupted source of power. The time to quench after a power loss is about 4 minutes



for our system. In principle, Yale has fast acting emergency backup power, but due to system upgrades and other issues, this emergency power was not available in early March 2016, when an unscheduled power outage resulted in a magnet quench, causing significant mechanical/structural damage to the DR and coaxial microwave lines. The repairs of this damage took two months, and fortunately none of the DR's  $^3\text{He}$  was lost; in spite of the mechanical damage, the gas flow lines were in no way compromised. Most of the forces due to this quench were caused by the large copper rings at the lower part of the gantry, and by the copper still shield extension. As noted above, we have replaced the copper still shield extension with another made from copper-plated stainless steel. In future designs, most of the gantry components will be constructed from copper-plated stainless steel, which should lead to more than a 100-fold reduction in forces during a quench.




Figure 3: Layout of the experiment. The red volume is the DR vacuum shield, the blue volume is the magnet body, and the orange shaded region is the volume bounded by the still shield and its extensions. The magnet's 70 K and 4 K shields are not shown below their interface with the fridge shields.

### 2.3. Microwave cavity

The initial microwave cavity for the experiment consists of a right circular cylinder of stainless steel, electroplated with oxygen-free high conductivity (OFHC) copper. After annealing, the cavity achieves a near-theoretical maximum value of the quality factor  $Q$ , as limited by the anomalous skin depth [9],

$$\delta_{\text{anom}} = \left( \frac{\sqrt{3} c^2 m_e v_F}{8\pi^2 \omega n e^2} \right)^{1/3}, \quad (6)$$

where  $m_e$  is the mass of the electron,  $n$  the electron density in the metal,  $v_F$  the Fermi velocity, and  $\omega$  the angular frequency. For copper,  $n = 8.5 \times 10^{22} \text{ cm}^{-3}$ , and  $v_F = 1.57 \times 10^8 \text{ m/s}$ .

The cavity employed for the initial data run was 25.4 cm in height and 10.2 cm in diameter. A single large diameter rod,  $\varnothing 5.1 \text{ cm}$ , was used to tune the cavity (see Fig. 2). The rod pivoted around an off-axis ceramic axle, such that its radial position could be adjusted from touching the wall to centered within the cavity. This produced a dynamic range of the  $\text{TM}_{010}$ -like mode of 3.6 – 5.8 GHz. To make finer frequency steps, we used a tuning vernier, a 3.2 mm diameter alumina rod that is attached to a tungsten weight. With the tuning rod centered, the cavity is more properly described as an annular cavity, which has the largest volume for a  $\text{TM}_{010}$ -like mode at a given frequency. The upper endcap had ports for two antennas, one with weak coupling used to measure the cavity's response in transmission and one whose insertion depth (and thus coupling  $\beta$ ) was variable. Around 5.8 GHz, typical values were  $C_{010} \sim 0.5$  for the form factor and  $Q_0 \sim 30,000$  for the unloaded cavity  $Q$ .

Our cavity development program has involved precision metrology and modeling in simulation environments such as CST Microwave Studio. Fig. 4 displays a simulation of the  $\text{TM}_{010}$ -like mode and a mode map of the cavity. Another innovation of this axion search is the use of the bead perturbation method to measure field profiles in a microwave cavity. To date, these have been axial profiles along the full length of the cavity at a single transverse position, derived by translating a small dielectric bead (alumina,  $\epsilon \sim 10$ ) on a Kevlar line controlled by a stepper motor and some pulleys. For a dielectric bead, the shift in mode frequency in perturbation theory is given by:

$$\frac{\Delta\omega}{\omega} = \frac{-(\epsilon - 1)}{2} \frac{V_{\text{bead}}}{V_{\text{cav}}} \frac{E(r)^2}{\langle E(r)^2 \rangle_{\text{cav}}^{1/2}} \quad (7)$$

where  $V_{\text{bead}}$  and  $V_{\text{cav}}$  indicate the volume of the dielectric bead and cavity, respectively [10]. For the  $\text{TM}_{010}$ -like mode, the electric field should be longitudinally invariant, and thus the shift in frequency as a function of axial distance should be constant. In the microwave cavity axion search, the bead perturbation technique is useful in the design and testing of new resonators, to determine the symmetry of a mode, to confirm or correct alignment

and machining tolerances to eliminate mode localization, and to identify mode-crossings regions where the modes have become substantially admixed, and thus the form factor  $C_{mnl}$  of the  $TM_{010}$ -like mode has begun to diminish (Fig. 5). A detailed study of microwave cavities for the dark matter axion experiment bringing together metrology, simulation, and EM characterization will be the subject of a future publication. For the present, a few observations in relation to experiments an order of magnitude lower in frequency [4, 7] may be useful.

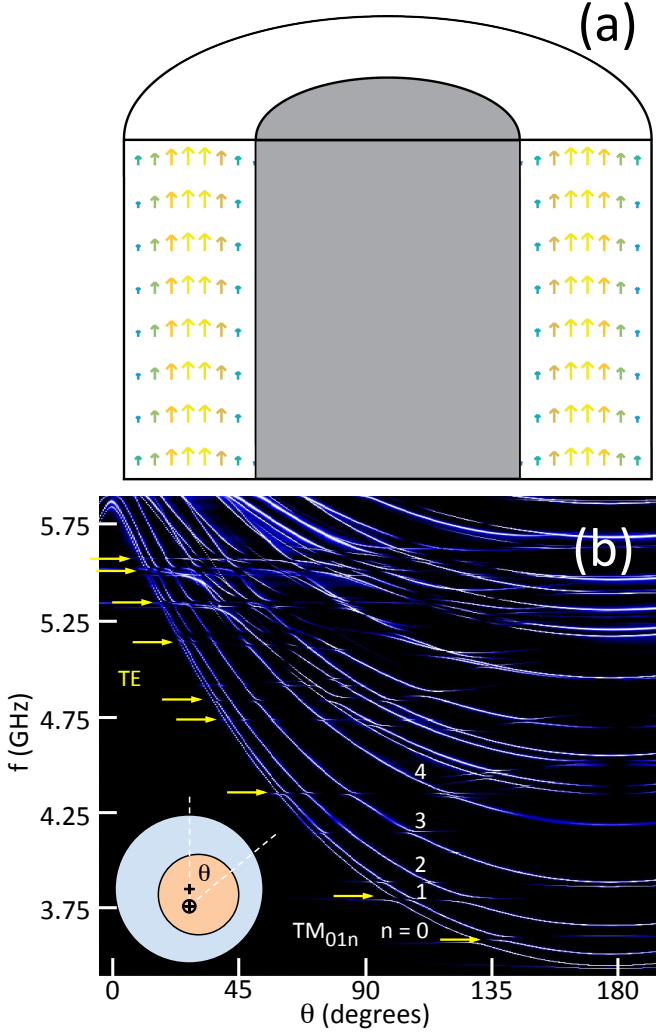


Figure 4: (a) CST Microwave Studio simulation of the  $TM_{010}$ -like mode with the tuning rod in the middle of the cavity. (b) Measured mode map of the cavity, as a function of the pivot angle of the tuning rod within the cavity. The frequencies of the  $TM_{01n}$  modes decrease steeply with increasing radial distance of the tuning rod from cavity center. TE modes (indicated by yellow arrows) become apparent approaching mode crossings with TM modes; the TE mode frequencies are largely insensitive to the position of the tuning rod.

First, for cavities of characteristic dimension  $O(10\text{ cm})$ , machining and alignment tolerances to avoid significant mode localization [11] are approaching the limit of good machine shop standards and readily available components.




Figure 5: Cavity spectra for two different rod positions, as the  $TM_{010}$ -like mode approaches a crossing with a high-order TE mode. For each rod position, the plots above the main panels show the shifts in frequency of the two modes, as a small dielectric bead is translated axially through the cavity. The frequency shift at each axial position  $Z$  is proportional to the square of the local electric field according to perturbation theory (see Eq. 7), thus elucidating the nature of the mode. (The spectra in the main panels are with the bead removed from the cavity.) (a) At a separation of 8 MHz, the purity of the  $TM_{010}$ -like mode is confirmed by the flat axial field profile shown on the top left. (b) At a separation of 2 MHz, the lower-frequency mode now clearly reflects a strong admixture of both a high-order TE mode along with the  $TM_{010}$ -like mode.

Even the transverse ‘play’ of the inner race of the bearing for the axle of the tuning rod, of order  $50 - 75\text{ }\mu\text{m}$ , can cause a tilt in the  $TM_{010}$  axial electric field profile of several percent. Second, modeling revealed that for the large tuning rod used in this experiment ( $r/R = 0.5$ ), the form factor  $C_{010}$  only approached the value associated with the idealized case (i.e. the rod electrically joined to the top and bottom endcaps) for gaps  $G < 250\text{ }\mu\text{m}$ . Third, it is becoming clear that engineered features necessary for practical cavities (e.g. penetrations for the antenna coupling, diagnostic ports, etc.) have a significant effect on the mode, and are likely to be more important at higher frequencies given the larger relative size of mechanical and electrical components penetrating the cavity.

#### 2.4. Mechanical Controls

The experiment has three mechanical systems: the rotary cavity tuning, the adjustment of the vernier, and the adjustment of the antenna. All mechanical controls are based on Applied Motion Products HT-23-595D NEMA 23, high-torque, double-shaft stepping motors, driven by model 5000-235 STR2 microstepping controllers. One side of each stepping motor double shaft is coupled to a ten turn potentiometer that encodes the net rotation angle. The vernier and antenna stepping motors are coupled to the experiment using 10:1 worm gear reductions, while the cavity tuning motor is coupled directly. The rotary motions are coupled into the vacuum using three K.J. Lesker FMH-25A Dynamic O-Ring Shaft Seal feedthroughs.

Within the cryostat, the antenna and vernier are controlled by 0.36 mm Kevlar thread lines that are wound directly onto the rotary feedthrough shaft that extends into the vacuum. At each stage of the DR, the lines pass through 3.2 mm ID, 2.5 cm long tubes that are thermally linked to the DR stage and serve as radiation shields. The lines are routed to the cavity by use of nylon pulleys with ball bearings (McMaster-Carr 3434T13) that have been ultrasonically cleaned to remove all lubricants that would freeze at low temperatures. It should be noted that Kevlar expands on cooling, and the spring tensioning of the lines must absorb the length change lest the lines loosen and possibly fall out of the pulley grooves.

The antenna and vernier are mounted on fixtures that allow motion along the cavity axis only. A spring is compressed when the antenna is pulled out of the cavity by the Kevlar thread, ensuring smooth and reversible motion. Gravity acting on the tungsten weight guarantees reversible motion for the vernier system.

Rotary motion for the tuning rod system is delivered to the mixing chamber level by use of a cryogenic G-10 tube (6.4 mm diameter, 0.79 mm wall). The ends of the G-10 tube are fitted with glued-in brass extensions which allow the use of set screws to couple to the feedthrough at the top, and to the mechanics at the mixing chamber end. At each DR stage, a 2.5 cm section of brass tubing (6.4 mm ID, 0.4 mm wall) is glued to the G-10, and a loose-fitting brass tube (1.9 cm long, 0.4 cm ID) is slipped over each 2.5 cm section and coupled to the stage using copper braid. At the 4 K stage, the G-10 tube inside has a brass blackbody radiation block.

Just below the mixing chamber, a pulley and torsion spring system is used to transfer rotary motion from the upper G-10 tube to a cryogenic G-10 tube (also 6.4 mm) along the DR and magnet axis, supported at the mixing chamber end by a ceramic bearing. A kevlar line connects the brass extension of the upper shaft to the 10.2 cm diameter pulley, and the torsion spring ensures that the line is always pulled taut. The lower G-10 shaft runs through the JPA magnetic shield and down to the cavity, where a 1.4:1 anti-backlash gear reduction provides the final radial displacement required to couple to the tuning rod axle.

These mechanical systems have generally provided the levels of control needed for the antenna and vernier; all three required appropriate selection of the microstepping resolution and had negligible heat loads at the operating temperature. The cavity bearings had some stiction due to alignment issues, and this resulted in sometimes erratic positioning, and a slow drift to final equilibrium after stepping. In operation of the experiment, the fine frequency control was done using the vernier, with less frequent larger rotations of the tuning rod which required up to 15 minutes of waiting for the rod to come to its equilibrium position. We plan to eliminate this problem by improving the cavity mechanics and replacing the room-temperature tuning rod drive system with an AttoCube piezoelectric rotator at base temperature.

## 2.5. Josephson Parametric Amplifier

A JPA is a nonlinear LC resonator capacitively coupled to a transmission line with an array of SQUIDs playing the role of the nonlinear inductance (see Fig. 6). By applying a DC flux through the SQUID loops, the resonant frequency can be tuned over several gigahertz. Driving the system near its resonance with a strong pump tone enables amplification of a weak signal nearby in frequency. The gain of this amplification can be traded off with the bandwidth over which it occurs by adjusting the pump power and frequency.




Figure 6: Microphotograph of JPA circuit. The SQUID array (approximately  $150\ \mu\text{m}$  long) is highlighted in red on the left; the circuit's resonance is determined by the SQUID inductance and the geometric capacitance (blue). The circuit is coupled to a  $50\ \Omega$  transmission line through a smaller capacitance (green). The surrounding superconducting ground plane is waffled in order to pin magnetic flux vortices in place and keep them from the SQUID array.

When the signal of interest is centered about the pump tone, the amplification process is noiseless (insofar as the JPA's internal loss is negligible) but phase-sensitive. In this mode of operation, the signal quadrature in phase with the pump is amplified by  $\sqrt{G}$ , while the one 90 degrees out of phase is squeezed by  $1/\sqrt{G}$ , where  $G$  is the single-quadrature power gain. When detuned completely to one side of the pump, the signal is amplified independent of its phase, but the JPA's intermodulation gain implies that an extra noise term enters from the image frequency, symmetric with the signal about the pump. In this configuration the added noise of a lossless JPA is equal to the thermal noise at the image frequency, giving rise to the quantum limit  $N_A = 1/2$  at zero temperature.

It is important to note that operating in the phase-sensitive mode by itself does not eliminate the half-photon of noise associated with the zero-point motion of the blackbody gas in the cavity (e.i., the second additive term in Eq. (4)). Moreover, the factor of two improvement in noise temperature for a lossless JPA in the phase-sensitive



mode is exactly canceled by the loss of information in one quadrature. Thus, phase-sensitive operation offers no improvement in axion search sensitivity unless we can also eliminate the zero-point motion of the input noise, e.g., by initializing the cavity in a squeezed state [12]. R&D with this aim is ongoing within the collaboration; for the present, we operate the JPA in the phase-insensitive mode.

The JPA currently installed in the experiment has a maximum resonant frequency of 6.5 GHz and a gain bandwidth product  $\sqrt{GB} = 31$  MHz. Half a flux quantum threaded through the area of a single SQUID loop in the array tunes the circuit through its full 2 GHz range; this extreme flux sensitivity necessitates specialized magnetic shielding, detailed in section 2.6. The circuit was fabricated using a Nb/Al-Ox/Nb trilayer process; the nonlinear inductance comes from 20 SQUIDs in series, where each SQUID comprises two Josephson junctions of critical current  $6 \mu\text{A}$  in parallel. The flux bias is delivered through a coil comprising about 20 turns of superconducting wire wound around the copper box housing the JPA chip.

The JPA can be turned into a passive mirror very simply by tuning off the pump tone and tuning the resonant frequency  $\sim 10$  linewidths from the frequency of interest via the flux bias; we use this process to calibrate the absolute JPA gain. We set a gain target of  $\simeq 21$  dB, below the onset of bifurcation inferred from deviations from a gaussian distribution in the JPA's quadrature noise spectra, but large enough to overwhelm the  $\sim 20$  quanta added noise of the HEMT amplifier which follows the JPA in the receiver chain. At this operating point the JPA bandwidth is about 2.8 MHz.




Figure 7: Schematic representation of the parameter space for JPA biasing. All features are intended for illustrative purposes only; they do not represent real measurements or calculations. The black curves are contours of constant gain, and the black arrows represent the path taken by the bias procedure outlined in the text, intersecting each gain curve at the minimum-power point. Beyond the critical point  $(\Delta_c, P_c)$ , the system begins to oscillate and can be bistable or multistable. We operate at lower power (red dot) to avoid this region of parameter space.

The procedure used to bias the JPA to this target gain

is illustrated schematically in Fig. 7. The detuning  $\Delta$  between the  $LC$  resonance and the pump frequency is adjusted by varying the flux bias. The gain is optimized with respect to detuning at constant pump power, and the pump power is increased if the gain remains too low after this optimization; the pump frequency remains fixed throughout the biasing process. With this procedure we always obtain within 0.2 dB of the target gain, and always operate at the point of lowest pump power for a given gain, where the JPA is most stable. The biasing procedure is fully automated and incorporated into the LabVIEW code that controls the data acquisition (see section 2.8). It typically takes  $\sim 6$  s to adjust the bias parameters after each cavity tuning step.

Maintaining high gain throughout the tuning range requires the pump power and flux bias to be controlled to 0.01 dB and 2 parts in  $10^5$ , respectively. The latter condition corresponds to 300 nA current resolution with our present bias coil, which we easily obtain using a 20-bit ADC with  $10 \mu\text{V}$  resolution and a homemade current source with 1 mA/V transconductance.

## 2.6. Magnetic Shielding

Shielding of the sensitive receiver components, in particular the JPA, from the stray field of the 9 T magnet presents a complicated engineering problem. Also of concern are magnetic gradients at the location of the JPA that lead to excess amplifier noise due to gain fluctuations in the presence of mechanical vibrations. We have successfully attained the degree of magnetic shielding required for stable JPA operation, as shown in Fig. 8(b).

The outer elements of the JPA shield can be seen in Fig. 8(a). The first level of shielding is from the bucking coils described in section 2.2. The next level of shielding comprises three persistent superconducting coils that provide a quasi-active shield against flux changes; given that they become superconducting before the 9 T field is ramped up, the flux inside will be maintained at the original ambient level. The axial spacing of the coils is not uniform because the bucked field gradient is asymmetric, with slightly higher flux at the lower end of the shield compared to the upper. The axial center of the shield is located at the field minimum.

The superconducting shield coils each comprise 100 turns of Supercon Inc. SC-T48B-M-0.7mm Cu-clad NbTi SC wire, with 0.43 mm diameter NbTi core (nominal), 0.70 mm diameter Cu cladding, and 0.75 mm diameter HML (Polyimide) insulation. The windings are on 9.3 cm ID brass hoop forms and potted with Stycast 2850 FT. After winding, the coil wire ends are trimmed to about 15 cm, and the last 5 cm of wire are annealed by heating to red heat using a MAPP gas torch; this also burns off the Polyimide insulation. The heated regions are cleaned and smoothed using 400 grit alumina emery paper. The wire ends were bonded [13] to make a persistent superconducting connection by use of a Koldweld KBM-9 wire coldwelding machine; the dies we had available were too



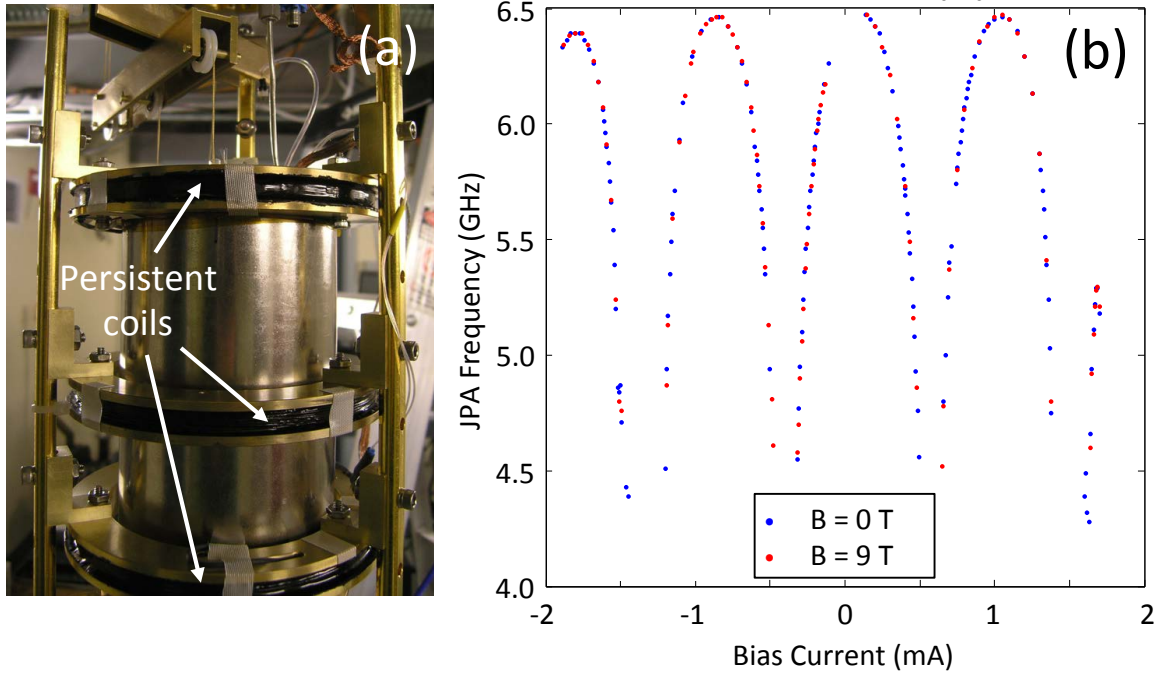


Figure 8: (a) Magnetic shielding of the JPA, including three superconducting persistent coils and (b) JPA frequency vs. flux bias current, showing no change in frequency as the main magnetic field is ramped from 0 – 9 T. The residual flux through the JPA is inferred to be  $\Delta\Phi < 0.01$  flux quantum.

large, so the outer copper cladding was built up through acid copper electroplating. After electroplating, the ends were trimmed back 1.6 mm, and the wire ends bonded. A micrograph analysis of an axially mechanically sectioned bond shows that the NbTi material flowed together at the center and pushed the copper uniformly in a radial direction from the center.

A similarly wound coil, with 24 cm OD, is clamped to the bottom of the DR mixing chamber plate. The fringe field of the 9 T magnet is about 300 G in this region; without the coil this is large enough to affect operation of the shielded cryogenic circulators (see section 2.7).

The ferromagnetic shield comprises a nested pair of cylinders made of 1.5 mm thick Amumetal 4 K, annealed per Amuneal’s proprietary process. The bottoms of the can have welded-on disks, while the tops are closed with tight-fitting lids that overlap 1.25 cm of the cylinder walls. The outer shield (FS1) has length 15.2 cm and OD 8.9 cm, while the inner shield (FS2) has length 13.3 cm and OD 7.6 cm. Both top lids have holes for the tuning rod (1.25 cm) and also for the JPA thermal link and transmission line. There are also three small holes, placed at 120° in both the top and bottom, that allow Kevlar tuning lines to pass through both shields.

Between the two ferromagnetic shields is a Pb superconducting shield, fabricated from 1.6 mm Pb sheet, which fits closely on the outside of FS2 and has a removable close-fitting lid. The Pb shield can is about 12.0 cm long, and covers FS2 to just below the lid. The Pb lid is about 3.8 cm long, giving an overlap with the cylinder of 2.5 cm. The

inside of the Pb can is coated with clear acrylic paint to avoid electrical contact and thermal currents, and is glued to FS2 using Loctite 680.

The final passive shield component is a 0.13 mm thick Nb sheet (99.8 % purity, Alfa-Aesar) that lines the inner surface of FS2. This sheet forms an “open cylinder” in that it has no endcaps, and there is an effective axial slit along its length which allows flux that would otherwise be trapped to escape. The sheet is long enough so that the ends overlap by about 1.25 cm, and the sheet is coated with acrylic paint, again to prevent the formation of electrical contacts. The sheet is glued to the inner surface using Loctite. This sheet enforces a boundary condition that the magnetic field must be axially homogeneous along its surface. The ferromagnetic boundary condition at the bottom and at the lid of FS2 is that the magnetic field is perpendicular to the surface, which matches perfectly the boundary condition imposed by the Nb sheet. The combined effect of the superconducting and ferromagnetic boundary conditions is to substantially reduce field gradients at the JPA.

A G-10 disk and nylon standoffs are used to support the JPA at the axial center of the shield, with a radial offset to allow the tuning rod to pass. The standoffs are screwed to 8-32 studs that extend from OFHC Cu spacers between the two ferromagnetic shields; the cooling of the shield assembly is through these studs and spacers which clamp the shields to each other and to a Cu plate affixed to the gantry. The JPA itself is thermally linked to the gantry via a Cu braid that exits through the top of the

FS2/Pb shield assembly.

A final active layer of magnetic shielding is provided by feedback to the JPA flux bias coil, discussed in section 2.7.

### 2.7. Receiver and microwave measurements

The cryogenic microwave layout and room-temperature microwave/IF layout are depicted schematically in Fig. 9 and Fig. A.13 in the appendix, respectively. The receiver signal path is shown in blue in both diagrams. This section will describe the microwave portion of the receiver, the RF input lines, and other detector functionality which is best understood in reference to Fig. 9 and Fig. A.13. The IF part of the receiver chain is discussed in section 2.8 below.

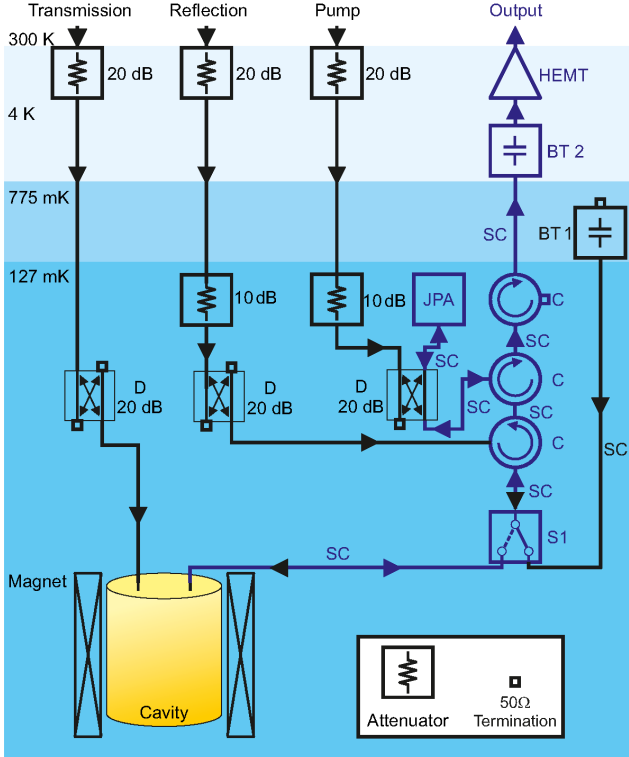


Figure 9: The cryogenic microwave layout. Blue arrows indicate the receiver signal path from the cavity to room temperature; black arrows indicate other paths used for network analysis, noise calibration, and JPA biasing. Component part numbers and manufacturers are listed in table A.1.

The input of the receiver is the cryogenic microwave switch S1, which may be toggled to connect the receiver to either the cavity or a termination thermally anchored to the DR still plate at  $T_H = 775$  mK. This switch allows us to calibrate the receiver’s added noise *in situ* via a procedure described in detail in section 2.9.

The JPA, playing the role of the preamplifier, comprises the magnetically shielded JPA circuit described in section 2.5, a directional coupler for the pump tone input, and a commercial ferrite circulator to separate input and output signals. Two other circulators are required to isolate the JPA from both the backaction of the second-stage amplifier and its own backaction in reflection from

the cavity. Signals exiting the JPA are amplified further by a HEMT amplifier at 4 K and another low-noise transistor amplifier at room temperature, and then routed both to a commercial vector network analyzer (VNA) and to an IQ mixer (M2) serving as the input of a homemade spectrum analyzer. The former path is used for cavity and receiver characterization, the latter for periodic noise calibrations and the cavity noise measurements constituting the axion search dataset.

Swept tones produced by the VNA may be directed via software-controlled switches S2 to any of three fridge input lines, through which they are transmitted through the cavity, reflected off the cavity, or sent directly to the JPA, bypassing the cavity. The first two lines are used to measure the cavity parameters  $\nu_0$ ,  $Q_L$ , and  $\beta$ ; the third line is used for JPA biasing and gain measurements, as discussed in section 2.5. Room-temperature attenuators were chosen to equalize sweep power incident on the JPA through each path. Attenuators at 4 K and base temperature reduce room-temperature thermal noise and the phase noise of the JPA pump generator to  $\sim$  mK contributions we can safely ignore.

0.085" NbTi/NbTi coaxial cables are used in the receiver signal path between base temperature and 4 K; these cables (marked SC in Fig. 9) remain superconducting in the 9 T field due to flux pinning. Stainless 0.085" coax is used between room temperature and 4 K in all four lines and down to base temperature in the input lines. All four coaxial lines are thermalized at each stage of the fridge with gold-plated copper clamps; the output line inner connectors are thermalized via a bias tee at 4 K with its DC input shorted to ground. NiCr resistors are used in all cryogenic attenuators and terminations.

The room-temperature microwave chain includes three signal generators from Keysight in addition to the VNA. An 8340B serves as the local oscillator (LO) for both the spectrum analyzer system and the flux feedback system discussed below. An E8257D with an ultra-low phase noise option ( $-120$  dBc/Hz at 100 kHz detuning from the carrier) provides the pump tone for the JPA. Finally an N5183B is used to inject synthetic axion-like signals into the cavity transmission line. All three generators, along with the VNA and the digitizer board used in the spectrum analyzer, lock to a common 10 MHz reference provided by a Stanford Research Systems (SRS) FS725m rubidium source (not pictured).

To generate synthetic axion signals with the expected linewidth ( $\sim 5$  kHz for a  $25 \mu\text{eV}$  axion), we inject band-limited white noise into the FM port of the N5183B; the linewidth is then controlled by the modulation depth. The synthetic axion power in the cavity is only known to  $\pm 3$  dB due to unknown cryogenic insertion losses of individual components. This uncertainty prevents us from using the synthetic axion system to precisely calibrate our sensitivity, but it is still useful as a way to validate both the data acquisition and analysis procedures. Our data acquisition software includes functionality for blind injection of ax-




Figure 10: (a) 100 Hz-resolution power spectra from individual 15-minute integrations around a frequency at which a synthetic axion signal was injected. (b) The optimally weighted combined spectrum, still at 100 Hz resolution. (c) The combined spectrum after rebinning to 5 kHz with weighting that takes into account the axion lineshape; the synthetic axion is clearly visible. See [14] for an overview of the analysis procedure, which will also be described more thoroughly in a forthcoming publication.

ion signals at random frequencies throughout a data run; Fig. 10 illustrates one such injection in our first run, with a nominal intracavity power level of  $-190$  dBm.

The room-temperature signal paths used exclusively by the JPA flux feedback system mentioned at the end of section 2.6 are shown in pink in Fig. A.13. The JPA flux bias is modulated at 26 Hz by adding an oscillating current (controlled by an Agilent 33220A function generator) to the DC bias current set in software. The modulation frequency is limited by eddy current shielding in the Cu JPA enclosure; the amplitude is set to yield a modulation depth of  $\sim 0.1$  flux quantum. When the DC flux bias is offset from the value that maximizes the JPA gain at fixed pump power and frequency, the JPA gain is modulated at the same frequency as the flux with phase determined by the sign of the offset. When the bias modulation is centered on the optimal value, only higher harmonics are present in the gain variation. Thus the JPA gain variation at the modulation frequency may be used as an error signal in an analog feedback system to stabilize the flux bias.

During noise measurements we use the VNA to inject a weak CW probe tone through the pump line at 30 kHz detuning from the pump. The function of the feedback circuitry in the upper left corner of Fig. A.13 is to provide an LO at the appropriate IF frequency to bring the probe tone to DC on the IF side of the mixers M3. Both quadrature outputs are amplified and squared using AD633 analog multipliers; the sum of squared signals is a measure of the received power in the probe tone and thus of the JPA

gain. A Stanford Research Systems SR510 lock-in amplifier measures the variation of the probe tone power at the modulation frequency. By adding the SR510 output to the bias current we obtain a simple proportional feedback loop; the SR510 gain, phase, and filtering are chosen to provide a stable feedback signal. During JPA biasing and sweep measurements the probe tone is not present because the VNA is otherwise occupied, so both modulation and feedback are interrupted by switching off the signal and reference outputs of the 33220A in software.

## 2.8. Data acquisition and operations

The spectrum analyzer comprises the discrete components between M2 and ADC in Fig. A.13, along with LabVIEW code that computes FFTs, implements image rejection, and averages power spectra in parallel with timestream data acquisition. During all noise measurements the LO for M2 is set 780 kHz above the  $\text{TM}_{010}$  mode frequency, the JPA pump is set 820 kHz below the mode, and both IF channels are sampled by the GaGe Oscar CSE4344 PCIe digitizer board at 25 MS/s.

Each IF channel is sensitive not only to the RF frequencies of interest in the lower sideband of the LO but also to unwanted image frequencies in the upper sideband. The  $90^\circ$  relative phase shift between the two otherwise identical IF outputs of an IQ mixer may be exploited for image rejection: adding the Q output to the I output with a  $+(-)90^\circ$  phase shift suppresses the upper (lower) LO sideband. This is the operating principle of commercial image reject mixers, which only work at particular IF frequencies because they implement the required  $90^\circ$  phase shift in hardware.

We implement image rejection in the frequency domain in software: this scheme works at any IF frequency, with the only possible drawback being that amplitude and phase mismatches in the discrete IF components in the two channels can limit the degree of image rejection. By taking FFTs of both the I and Q channels and defining  $X(\omega) = (\text{Re}[I(\omega)] - \text{Im}[Q(\omega)] + i(\text{Im}[I(\omega)] + \text{Re}[Q(\omega)]))$  we obtain rejection of the upper sideband better than 20 dB at all IF frequencies of interest in the power spectrum  $|X(\omega)|^2$ .

More specifically, 14-bit ADCs on the GaGe board digitize both channels simultaneously in 5 s segments, then transfer each segment to PC RAM. The total data in each segment across both IF channels is 437.5 MB, and the time required to transfer this data to RAM is 1.2 s, which caps the data acquisition efficiency at 80%. The *in situ* processing code divides each segment into 500 non-overlapping 10 ms records in each channel, computes the FFT of each record with no windowing, combines the I and Q FFTs corresponding to the same 10 ms time slice to implement image rejection, constructs a power spectrum from each sample of  $X(\omega)$ , and averages all 500 power spectra corresponding to each segment. All processing for each 5 s segment occurs in RAM in parallel with the acquisition of

the next segment. At the end of the data acquisition period (typically 15 minutes), the power spectra from all segments are averaged together to obtain a single spectrum obtained from  $100 \text{ Hz} \times 15 \text{ minutes} = 9 \times 10^4$  averages, which is written to disk.

Thus the output at each tuning step is a heavily averaged power spectrum with 100 Hz resolution extending from DC to the 12.5 MHz Nyquist frequency. The usable IF bandwidth extends to  $\sim 2.5$  MHz, limited by the bandwidth of the low-pass filters F2. The sampling rate is set so far above this to eliminate the need for a high-order filter with significant passband ripple on the output of the IF amps A2. The function of the F1 filters is to attenuate RF leakage into the IF chain.

The 100 Hz spectral resolution is much smaller than the  $\sim 5$  kHz expected linewidth of a virialized axion signal – thus data from this detector can be used to set more stringent limits on the abundance of non-virialized axions with  $Q_a \lesssim 10^7$  without any additional hardware. A technical advantage of the narrow resolution is that it enables us to better identify and eliminate individual IF bins contaminated by narrowband interference. Empirically some of these “IF spikes” are associated with ground loops and others are due to room-temperature electronics. The DC blocks and baluns in Fig. A.13 were added in commissioning to eliminate as many of the IF spikes as possible; remaining spikes are flagged and removed as part of the analysis procedure.




Figure 11: Diagram illustrating the IF setup for the experiment. Note that the IF frequency axis is reversed relative to the RF frequencies of these features. The JPA gain profile and  $\text{TM}_{010}$  Lorentzian profile are plotted using real data and a fit to real data, respectively. Both plots have logarithmic y axes; the absolute scale of the sensitivity plot is arbitrary. The red dot-dashed arrow indicates the probe tone created on the opposite side of the pump by the JPA’s intermodulation gain. Images of the axion-sensitive Fourier components around the cavity, created on the other side of the pump by the same process, are omitted in the diagram for clarity.

Fig. 11 is a simplified diagram of the IF setup showing the most relevant features discussed above, and illustrating that we further limit our analysis to 1.302 MHz centered on the cavity. The width of this analysis band is roughly twice the maximum cavity linewidth; axion conversion at larger detunings from the cavity mode frequency contributes negligibly to the SNR. The LO frequency must lie outside the analysis band to avoid superimposing different analysis band Fourier components, and the pump

frequency must lie outside the analysis band for the JPA to operate in the desired phase-insensitive mode (see section 2.5). The LO and pump frequencies also cannot be equal, or else the I and Q channels will pick out (not necessarily equal) linear combinations of the JPA’s amplified and squeezed quadratures, and image rejection will fail.

These constraints indicate that the analysis band should be centered roughly halfway between the LO and pump in each spectrum. If the analysis band were too close to the pump, the spectrum would be contaminated by both the feedback probe tone and the pump tone’s phase noise. If the analysis band were too close to DC,  $1/f$  noise would dominate, and the relative contribution of the 4 K HEMT amplifier to the receiver’s added noise would grow with detuning from the pump frequency. The precise positioning of the analysis band was tweaked to exclude bins in which IF interference was most persistent.

Data acquisition for the experiment is fully automated and controlled by a LabVIEW program. At the beginning of each iteration, the DAQ program tunes the  $\text{TM}_{010}$  mode, measures  $\nu_0$  and  $Q_L$  with a VNA sweep through the cavity transmission line, then sets the LO and pump frequencies based on the new value of  $\nu_0$ . It optimizes the JPA gain as described in section 2.5, measures the gain profile over 5 MHz centered on the pump, turns on the flux feedback system, and makes the 15-minute power spectrum measurement described above. At the end of the power spectrum measurement, the JPA gain and cavity transmission are each measured again, to identify drifts of the cavity mode associated with mechanical relaxation (see section 2.3) and unusually large bias flux jumps that the feedback system was unable to correct for; roughly 0.6% of spectra were cut from our initial analysis due to anomalous flux or frequency drifts. Finally,  $\beta$  is measured with a reflection sweep. Noise calibrations (see section 2.9) are also performed intermittently to constrain variation of  $T_S$  throughout the run. The average live time efficiency during the run is 72%.

All sweep data is saved to disk along with the 100 Hz averaged power spectrum and critical parameter values such as the LO frequency; this amounts to about 3 MB per iteration. Data is stored locally at Yale and also transferred to a remote server at Berkeley for back-up and long term storage. The offline analysis procedure is outside of the scope of this paper, and will be the subject of a forthcoming publication.

## 2.9. Noise calibration

A simple way to measure the added noise of any microwave receiver is via a Y-factor measurement, in which we connect the receiver input to two  $50 \Omega$  loads at known temperatures  $T_C$  and  $T_H$ , and measure the hot/cold noise power ratio

$$Y = \frac{P_H}{P_C} = \frac{N_H + N_A}{N_C + N_A}, \quad (8)$$



where  $N_C(N_H)$  is the thermal noise from the cold (hot) load including the zero-point contribution, and  $N_A$  is the receiver's added noise. We can then solve Eq. (8) for  $N_A$  to obtain

$$N_A = \frac{N_H - Y N_C}{Y - 1}. \quad (9)$$

In our experiment,  $N_H$  comes from a 775 mK termination, and  $N_C$  comes from the cavity (on resonance) or from a terminated port on the directional coupler in the reflection input line, reflected off the cavity (off resonance; see Fig. 9).  $N_A$  includes the added noise of the JPA preamplifier and also the added noise of subsequent amplifiers referred to the JPA input; of those, in practice only the 4 K HEMT contribution is non-negligible.

Eq. (3) indicates that the effect of imperfect power transmission efficiency  $\eta$  between the cavity and the JPA is to rescale the total noise by  $1/\eta$ . It is convenient to define  $\eta = \eta_0 \eta_1$ , with  $\eta_0$  between the cavity and S1 and  $\eta_1$  between S1 and the JPA input. Then  $N_A$  measured by the Y-factor method will also include an additional additive term

$$N_\eta = \frac{1 - \eta_1}{\eta_1} (N_C + N'_A), \quad (10)$$

where  $N'_A$  indicates the sum of amplifier contributions. This term arises because the lossy elements that contribute to  $\eta_1$  are at temperature  $T_C$  (see Fig. 9), and thus attenuate both the hot load noise and the axion signal but not the cold load noise. Thus our measurement of  $N_A$  yields only the net effect of the JPA added noise, HEMT noise, and  $\eta_1$ , not any of these components individually. The effect of  $\eta_0$  on the SNR cannot be measured *in situ*; we estimate 0.6 dB from the long superconducting cable and connector losses, and take this loss into account in calculating our exclusion limits.

The general expression in Eq. (9) also must be modified to account for several specific features of our receiver. First, as noted in section 2.5, in our present configuration the JPA's added noise is simply the thermal noise on the opposite side of the pump from the cavity, coupled into the analysis band via image gain. Thus, the JPA contribution to  $N_A$  differs (by a known amount) in the two switch configurations. Second, both sides of Eq. (9) are in principle functions of IF frequency within any given measurement; in particular the lower JPA gain at the low-IF-frequency end of the analysis band (see Fig. 11) implies less suppression of the HEMT noise and thus larger  $N_A$ . We take both of these effects into account by generalizing Eq. (9) to

$$N_A(\nu) = \frac{2[A(\nu)N_H - Y(\nu)N_C]}{Y(\nu) - A(\nu)} + N_C, \quad (11)$$

where  $A(\nu) = G_H(\nu)/G_C(\nu)$  is obtained from the measured hot/cold JPA gain profiles.

During commissioning of the experiment, we consistently obtained  $N_A \simeq 1.35$  quanta from Y-factor measurements far detuned from the cavity mode independent of




Figure 12: A representative noise measurement.  $N_C$  (black dotted line) is obtained from thermometry,  $N_A$  (red dot-dashed line) is obtained from the average of off-resonance Y-factor measurements, and  $N_{\text{cav}}$  (blue dashed line) is from a single Y-factor measurement during the data run.  $N_{\text{sys}}$  (pink solid line) is the sum of these contributions.

RF frequency. This result is consistent with the added noise of an ideal nondegenerate JPA at  $T_C$  (0.63 quanta),  $\simeq 0.2$  quanta input-referred HEMT noise, and an additional  $\simeq 0.5$  quanta which we can attribute to  $\simeq 2$  dB loss before the JPA. This is a plausible value result for  $\eta_1$ , if a little on the high side.

In cold load noise measurements near the  $\text{TM}_{010}$  frequency we observe an additional Lorentzian excess  $N_{\text{cav}}(\nu)$  centered on the IF frequency of the mode. Naively applying Eq. (11) to such measurements would treat the  $N_{\text{cav}}$  as a contribution to the added noise (because this expression assumes the input noise in the cold load is completely characterized by the single number  $T_C$ ), and therefore overestimate the effect on the total noise  $N_{\text{sys}}$ , since  $N_A$  appears in both the numerator and denominator of Eq. (8), whereas  $N_C$  appears only in the denominator.

Therefore, we make the ansatz that the added noise is the same on-resonance as off-resonance, which allows us to quantify the peak excess contribution  $N_{\text{cav}}(\nu_0) \simeq 1$  quantum. During detector commissioning we examined the dependence of  $N_{\text{cav}}$  on the mixing chamber plate temperature  $T_C$ , and observed that the excess vanished for  $T_C \gtrsim 550$  mK. This behavior indicates that  $N_{\text{cav}}$  has a thermal origin as opposed to being an artifact of feedback between the JPA and cavity mode. Studies of the dependence of  $N_{\text{cav}}$  on JPA gain and on  $\beta$  further disfavor any sort of feedback explanation, which validates the above ansatz.

Having established that  $N_{\text{cav}}$  is of thermal origin, the most likely culprit is a poor thermal link between the tuning rod and the cavity barrel, as the only thermal connection in the present design is through a thin-walled alumina shaft. An improved thermal link through the axle of the

tuning rod has been designed and is being implemented.

In practice, our *in situ* noise calibration procedure consists of the following steps, repeated for each switch configuration: we bias up the JPA to target gain, measure the gain profile, and take 5 s of noise data, from which we construct a power spectrum with 10 kHz resolution, for a total of 130 data points in the analysis band. We wait 2 minutes for thermal transients to die down after actuating the switch in either direction; the heat due to each switch actuation is 10 mJ.

From these measurements we obtain a typical total noise of  $N_{\text{sys}} \simeq 3$  quanta on resonance, falling to  $\simeq 2.2$  quanta at the edges of the analysis band. The overall fractional error  $\delta N_{\text{sys}} \sim 6\%$  is dominated by  $\sim 17\%$  uncertainty in  $N_{\text{cav}}$ ;  $\delta N_A$  is only  $\sim 4\%$ , and the change in  $N_C$  is negligible even allowing for a  $\pm 20$  mK calibration error in the mixing chamber thermometry. Individual contributions to  $N_{\text{sys}}$  are plotted in Fig. 12 for a representative measurement. This is the lowest noise demonstrated to date in a microwave cavity axion search.

### 3. Summary and Conclusions

The experiment achieved its design sensitivity of  $g_\gamma = 2.3 \cdot g_\gamma^{KSVZ}$  in its first run, in spite of the imperfect thermal link between the tuning rod and the rest of the cavity. The improved thermal link should bring the tuning rod into thermal equilibrium with the mixing chamber and improve the linear coupling sensitivity by 20%.

That sensitivity this close to the KSVZ line can be reached in an experiment of 1.5 L volume (c.f. the 200 L volume of the ADMX detector [4, 7]) underscores the role of technology in the microwave cavity experiment. This platform is now poised to explore innovative concepts in both amplifiers and cavities to increase sensitivity and decrease scan time. A receiver based on injecting a squeezed state of vacuum into the cavity by one JPA and reading it out with another will be the first such innovation to be explored [12]. Beating the Standard Quantum Limit with such a receiver has been demonstrated on the bench [15], but this will be the opportunity to validate the feasibility of such a receiver in the rigors of an actual operating environment.

On the cavity side, a photonic band gap resonator is being designed to eliminate all interfering TE modes. As the TM mode of interest is tuned in frequency, avoided crossings with TE modes block a significant fraction of frequency coverage. This loss of coverage has been the bane of rapid and efficient mass coverage to date. Such structures have been thoroughly explored in accelerator physics [16, 17]; their successful adaptation to the microwave cavity axion experiment will largely pivot on an effective and reliable means of tunability. Similarly, cavities are being explored adapting design principles of distributed Bragg reflector structures, which attain much higher quality factors  $Q$ , providing greater sensitivity, and thereby scan rate, for the experiment [18].

### 4. Acknowledgements

This work was supported under the auspices of the National Science Foundation, under grants PHY-1067242, and PHY-1306729, the Heising-Simons Foundation under grants 2014-181, 2014-182, and 2014-183, and the U.S. Department of Energy by Lawrence Livermore National Security, LLC, Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344. We thankfully acknowledge the critical contributions by Matthias Bühler of Low Temperature Solutions UG for design of and upgrades to the cryogenic system. MS is supported by the National Science Foundation Graduate Research Fellowship Program under Grant No. DGE-1106400.

### References

- [1] P. Graham, I. Irastorza, S. Lamoreaux, A. Lindner, K. A. van Bibber, Experimental Searches for the Axion and Axion-Like Particles, *Annual Review of Nuclear and Particle Science* 65 (2015) 485–514. doi:10.1146/annurev-nucl-102014-022120.
- [2] P. Sikivie, Experimental Tests of the “Invisible” Axion, *Phys. Rev. Lett.* 51 (1983) 1415. doi:10.1103/PhysRevLett.51.1415.
- [3] P. Sikivie, Detection rates for “invisible”-axion searches, *Phys. Rev. D* 32 (1985) 2988–2991. doi:10.1103/PhysRevD.32.2988.
- [4] H. Peng, S. Asztalos, E. Daw, N. Golubev, C. Hagmann, D. Kinion, J. LaVeigne, D. Moltz, F. Nezrick, J. Powell, L. Rosenberg, P. Sikivie, W. Stoeffl, N. Sullivan, D. Tanner, M. Turner, K. van Bibber, Cryogenic cavity detector for a large-scale cold dark-matter axion search, *Nucl. Instrum. Meth. A* 444 (2000) 569–583.
- [5] R. H. Dicke, The measurement of thermal radiation at microwave frequencies, *Review of Scientific Instruments* 17 (7) (1946) 268–275.
- [6] C. M. Caves, Quantum limits on noise in linear amplifiers, *Phys. Rev. D* 26 (8) (1982) 1817.
- [7] S. Asztalos, G. Carosi, C. Hagmann, D. Kinion, K. van Bibber, M. Hotz, L. Rosenberg, G. Rybka, A. Wagner, J. Hoskins, C. Martin, N. Sullivan, D. Tanner, R. Bradley, J. Clarke, Design and performance of the ADMX SQUID-based microwave receiver, *Nucl. Instrum. Meth. A* 656 (2011) 39–44. doi:10.1016/j.nima.2011.07.019.
- [8] D. Tanner, private communication (2011).
- [9] C. Kittel, *Introduction to Solid State Physics*, 8th Edition, John Wiley & Sons, Inc., New York, 2005.
- [10] J. C. Slater, Microwave Electronics, *Rev. Mod. Phys.* 18 (1946) 441–512. doi:10.1103/RevModPhys.18.441. URL <http://link.aps.org/doi/10.1103/RevModPhys.18.441>
- [11] C. Hagmann, P. Sikivie, N. Sullivan, D. Tanner, S. Choi, Cavity design for a cosmic axion detector, *Rev. Sci. Instrum.* 61 (3) (1990) 1076–1085.
- [12] H. Zheng, M. Silveri, R. T. Brierley, S. M. Girvin, K. W. Lehnert, Accelerating dark-matter axion searches with quantum measurement technology (Jul. 2016). arXiv:1607.02529.
- [13] G. Luderer, P. Dullenkopf, G. Laukien, Superconducting joint between multifilamentary wires, *Cryogenics* 14 (9) (1974) 518–519.
- [14] B. M. Brubaker, et al., First results from a microwave cavity axion search at 24  $\mu\text{eV}$  (submitted to *Physical Review Letters*) (2016). arXiv:1610.02580.
- [15] F. Mallet, M. Castellanos-Beltran, H. Ku, S. Glancy, E. Knill, K. Irwin, G. Hilton, L. Vale, K. Lehnert, Quantum state tomography of an itinerant squeezed microwave field, *Phys. Rev. Lett.* 106 (2011) 220–202. doi:10.1103/PhysRevLett.106.220502.

- [16] E. Smirnova, C. Chen, M. Shapiro, J. Sirigiri, R. Temkin, Simulation of photonic band gaps in metal rod lattices for microwave applications, *J. Appl. Phys.* 91 (2002) 960–968. doi:10.1063/1.1426247.
- [17] E. Nanni, S. Lewis, M. Shapiro, R. Griffin, R. Temkin, Photonic-Band-Gap Traveling-Wave Gyrotron Amplifier, *Phys. Rev. Lett.* 111 (2013) 235101. doi:10.1103/PhysRevLett.111.235101.
- [18] C. Flory, R. Taber, High Performance Distributed Bragg Reflector Microwave Resonator, *IEEE Transactions on Ultrasonics, Ferroelectrics and Frequency Control* 44 (1997) 486–495. doi:10.1109/58.585133.
- [19] G. E. Theriault, A. J. Leidich, T. H. Campbell, Application of the CA3018 Integrated-Circuit Transistor Array, Tech. Rep. AN5296.0, Intersil (2012).

## Appendix A. Room-Temperature Receiver Layout

Table A.1: Component part numbers.

	Label	Type	Supplier: Part #
Cryogenic	BT1	Bias tee	Mini-Circuits: ZX85-12G+ (ferrite removed)
	BT2	Bias tee	Anritsu: K250
	C	Circulator	Pamtech Inc.: CTH1184K18
	D	Directional coupler	Pasternack: PE2211-20
	HEMT	HEMT amplifier	Low Noise Factory: LNF-LNC4.8A
	S1	Switch	Radiall: R577443005
	SC	NbTi/NbTi coax	Keycom: UPJ07
Room-temperature	A1	RF amplifier	Miteq: AMF-4F-04001200-15-10P
	A2	IF amplifier	Homemade: based on Fig. 2 in [19]
	A3	IF amplifier	Mini-Circuits: ZFL-500LN-BNC+
	A4	IF amplifier	Stanford Research: SR445A
	A5	IF amplifier	Stanford Research: SR560
	AT1	Step attenuator	Agilent Technologies: 8496H
	AT2	Step attenuator	Agilent Technologies: 8494H
	B	Balun	North Hills: 0017CC
	DC	DC block	Inmet: 8039
	F1	Low-pass filter	Mini-Circuits: VLFX-80
	F2	Low-pass filter	Mini-Circuits: SLP-1.9+
	F3	Low-pass filter	Mini-Circuits: BLP-2.5+
	F4	Low-pass filter	Homemade: 3.39 kHz single pole RC filter
	I1	Isolator	Ditom Microwave: D314080
	I2	Double isolator	Ditom Microwave: D414080
	M1	Mixer	Marki Microwave: M1-0408
	M2	IQ mixer	Marki Microwave: IQ0307LXP
	M3	Mixer	Mini-Circuits: ZAD-8+
	PD1	Power splitter/combiner	Mini-Circuits: ZX10-2-71-s+
	PD2	Power splitter/combiner	Mini-Circuits: ZFRSC-2050+
	S2	Switch	Mini-Circuits: ZFSWA2-63DR+



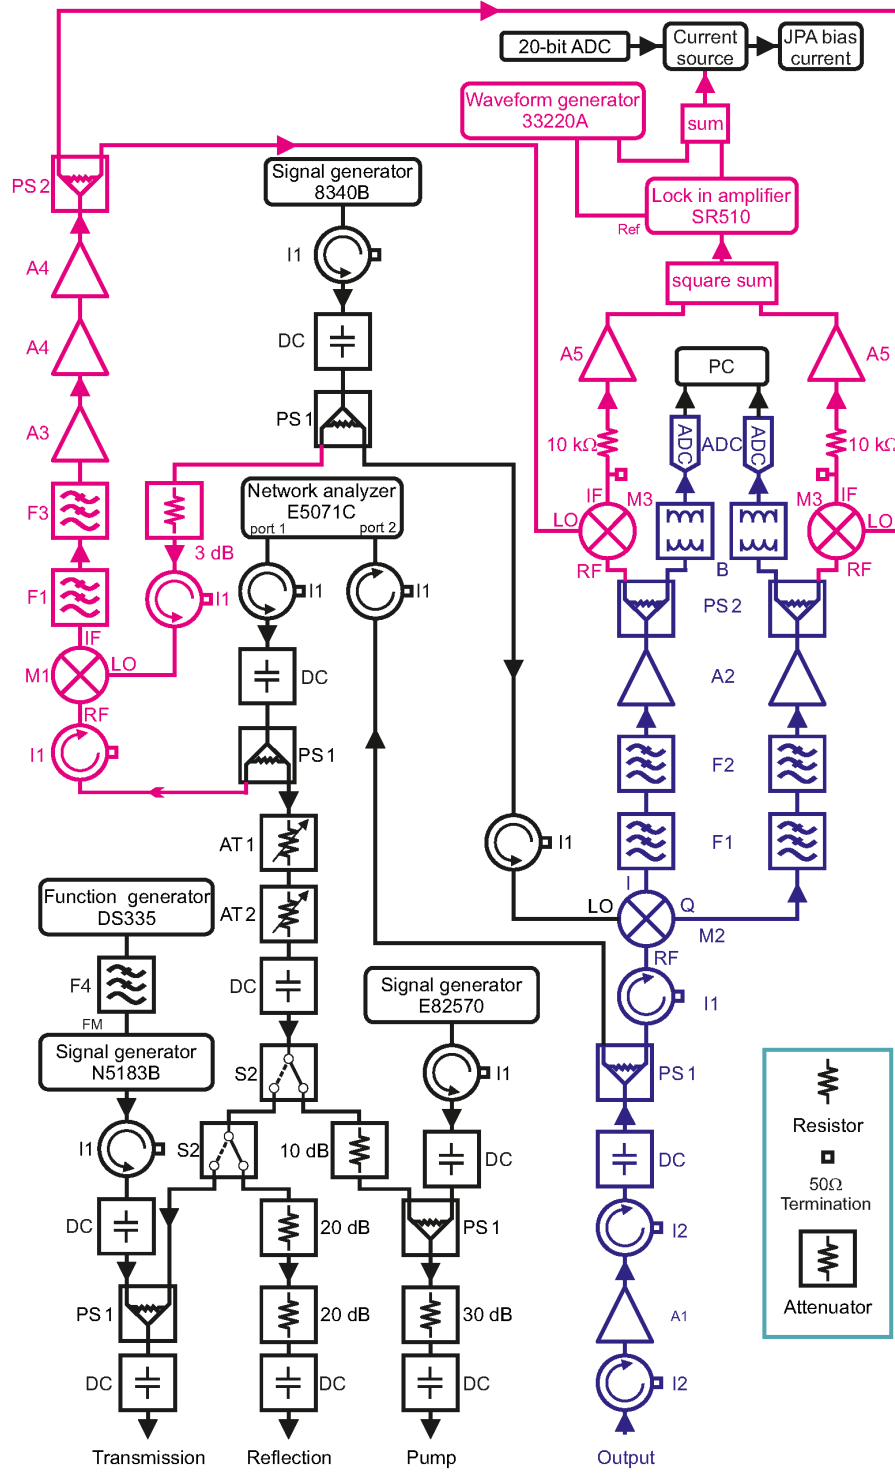


Figure A.13: The room-temperature microwave/IF layout. Blue arrows indicate the receiver signal path from the top of the fridge to the ADCs; black arrows indicate other paths used for network analysis, JPA biasing, providing LO power, and synthetic axion signal injection. Parts of the chain used exclusively by the JPA flux feedback system are indicated in pink. Component part numbers and manufacturers are listed in table A.1; those shown on the diagram are from Keysight or Stanford Research Systems.